arXiv:1703.10476v1 [cs.CV] 30 Mar 2017

Speaking the Same Language: Matching Machine to Human Captions by Adversarial Training

Rakshith Shetty[1] Marcus Rohrbach[2] Lisa Anne Hendricks[2]
Mario Fritz[1] Bernt Schiele[1]

[1]Max Planck Institute for Informatics, Saarland Informatics Campus, Saarbrücken, Germany
[2]UC Berkeley EECS, CA, United States

Abstract

While strong progress has been made in image captioning over the last years, machine and human captions are still quite distinct. A closer look reveals that this is due to the deficiencies in the generated word distribution, vocabulary size, and strong bias in the generators towards frequent captions. Furthermore, humans – rightfully so – generate multiple, diverse captions, due to the inherent ambiguity in the captioning task which is not considered in today's systems.

To address these challenges, we change the training objective of the caption generator from reproducing ground-truth captions to generating a set of captions that is indistinguishable from human generated captions. Instead of handcrafting such a learning target, we employ adversarial training in combination with an approximate Gumbel sampler to implicitly match the generated distribution to the human one. While our method achieves comparable performance to the state-of-the-art in terms of the correctness of the captions, we generate a set of diverse captions, that are significantly less biased and match the word statistics better in several aspects.

Ours: a person on skis jumping over a ramp

Ours: a skier is making a turn on a course

Ours: a cross country skier makes his way through the snow

Ours: a skier is headed down a steep slope

Baseline: a man riding skis down a snow covered slope

Figure 1: Four images from the test set, all related to skiing, shown with captions from our adversarial model and a baseline. Baseline model describes all four images with one generic caption, whereas our model produces diverse and more image specific captions.

1. Introduction

Image captioning systems have a variety of applications ranging from media retrieval and tagging to assistance for the visually impaired. In particular, models which combine state-of-the-art image representations based on deep convolutional networks and deep recurrent language models have led to ever increasing performance on evaluation metrics such as CIDEr [33] and METEOR [7] as can be seen e.g. on the COCO image Caption challenge leaderboard [6].

Despite these advances, it is often easy for humans to differentiate between machine and human captions – in particular when observing multiple captions for a single image. As we analyze in this paper, this is likely due to artifacts and deficiencies in the statistics of the generated captions, which in turn becomes more apparent when multiple samples are observed. More specifically, we observe that state-of-the-art systems frequently "reveal themselves" by generating a different word distribution and using smaller vocabulary. An even closer look shows that generalization from the training

Ours	a bus that has pulled into the side of the street a bus is parked at the side of the road a white bus is parked near a curb with people walking by	a group of people standing outside in a old museum an airplane show where people stand around a line of planes parked at an airport show
Base line	* a bus is parked on the side of the road * a bus that is parked in the street a bus is parked in the street next to a bus	a group of people standing around a plane a group of people standing around a plane a group of people standing around a plane

Figure 2: Two examples images with multiple captions generated by our model compared to a baseline model. Bi-grams which are top-20 frequent bi-grams in the training set are highlighted in red (e.g., "a group" and "group of"). Captions which are replicas from training set are marked with an asterisk.

set is still challenging and generation is biased to frequent fragments and captions.

Also, today's systems are evaluated to produce a single caption. Yet, multiple potentially distinct captions are typically correct for a single image – a property that is reflected in human ground-truth. This diversity is equally not reproduced by state-of-the-art caption generators [34, 20].

Therefore, our goal is to make image captions less distinguishable from human ones – an approach that is in the spirit of a Turing Test. We also embrace the ambiguity of the task and extend our investigation to predicting sets of captions and evaluating their quality in particular in terms of the diversity in the generated set. In contrast, predominant approaches to image captioning are trained with an objective to reproduce the captions as provided by the ground-truth.

Instead of relying on handcrafting different loss-functions to achieve our goal, we propose an adversarial training mechanism for image captioning. For this we build on Generative Adversarial Networks (GANs) [13], which have previously successfully mainly been used to generate continuous data distributions such as generating images [8, 27], although exceptions exist [23]. In contrast to images, captions are discrete, which poses a challenge when trying to backpropagate through the generation step. To overcome this obstacle, we use a Gumbel sampler [17, 24] that allows for end-to-end training.

In detail, we look at the problem for caption set prediction for images and discuss metrics to measure the caption diversity and compare it to human ground-truth. We contribute a novel approach towards this problem using an adversarial formulation. The evaluation of our model shows that overall system accuracy is on par to the state-of-the-art, but we greatly increase the diversity of the caption sets and closely match the ground-truth statistics in several measures. Our model produces qualitatively more diverse captions both across images containing similar content (see Figure 1) and when sampling multiple captions for each image (see Figure 2).

2. Related Work

Image Description. Early visual description models rely on first recognizing visual elements, such as objects, attributes, and activities, and then generating a sentence using a variety of methods such as a template model [12], n-gram language model [19], or statistical machine translation [29]. Advances in deep learning have led to end-to-end trainable models which combine deep convolutional networks to extract rich visual descriptors and recurrent networks to generate sentences [10, 35, 18, 25].

Though modern description models are capable of producing coherent sentences which accurately describe an image, they tend to produce generic sentences which are replicated from the train set [9]. Furthermore, an image can correspond to many valid descriptions. However, at test time, sentences generated with methods such as beam search are generally very similar. [34, 20] focus on increasing sentence diversity by integrating a diversity promoting heuristic into standard beam search models. [36] attempts to increase the diversity in caption generation by training an ensemble of caption generators each specializing in different portions of the training set. In contrast, we focus on improving diversity of generated captions using a single model. Our method achieves this by learning a corresponding model using a different training loss as opposed to after training has completed. We note that generating diverse sentence is also a challenge in language-only dialogue generation and has been studied in the linguistic community, see e.g. [20, 21].

When training recurrent description models, the most common method is to predict a word w_t conditioned on an image and all previous *ground truth* words. At test time, each word is predicted conditioned on an image and previously *predicted* words. Consequently, at test time predicted words may be conditioned on words that were incorrectly predicted by the model. By only training on ground truth words, the model cannot effectively learn to recover when it predicts an incorrect word during training. To avoid this, [4] propose a scheduled sampling training scheme in which begins by training with ground truth words, but then slowly

conditions generated words on words previously produced by the model. When we train our model, we also condition on predicted words which allows our model to learn how to effectively recover from mistakes during generation.

In this work, we formulate our sentence generator as a generative adversarial network. In addition to an image input, we also input a noise vector and train the generator with an adversarial loss. Our discriminator model explicitly encourages generated sentences to be diverse and indistinguishable from human generated sentences. Consequently, our model generates sentences the better reflect the way humans describe images while maintaining similar correctness as determined by a human evaluation.

Generative Adversarial Networks. The Generative Adversarial Networks (GANs) [13] framework learns generative models without explicitly defining a loss from a target distribution. Instead, GANs learn a generator by receiving a loss from a discriminator which tries to differentiate between real and generated samples, where the generated samples come from the generator. When training to generate real images, GANs have shown encouraging results [8, 27]. In all these works the target distribution is continuous. In contrast our target, a sequence of words, is discrete. Applying GANs to discrete sequences is challenging as it is unclear how to best back-propagate the loss through the sampling mechanism.

A few works have looked at generating discrete distributions using GANs. [23] aim to generate a semantic image segmentation with discrete semantic labels at each pixel. Most similar to our work is the concurrent work [21] who also generate text with a GAN, however, they aim to generate a diverse language dialogue of multiple sentences while we aim to diverse sentences for a single image. While [21] rely on the reinforcement rule [37] to handle back-propagation through the discrete distribution, we use the Gumbel Softmax [17], see Section 3.1 for further discussion on the difference. The generator training in [21] concurrently uses both the adversarial loss and the maximum likelihood update for each training step. We however train the generator with purely adversarial loss after pre-training.

3. Adversarial Caption Generator

The image captioning task can be formulated as follows: given an input image x the generator G produces a caption, $G(x) = [w_0, \ldots, w_{n-1}]$, describing the contents of the image. There is an inherent ambiguity in the task, with multiple possible correct captions for an image, which is also reflected in diverse captions written by human annotators (we quantify this in Table 4). However, most image captioning architectures ignore this diversity during training. The standard approach to model $G(x)$ is to use a recurrent language model conditioned on the input image x [10, 35], and train it using a maximum likelihood (ML) loss considering every

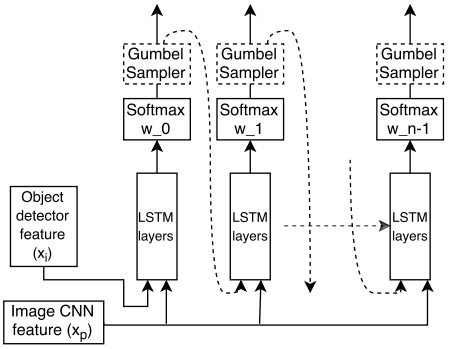

Figure 3: Block diagram of the caption generator

image–caption pair as an independent sample. This ignores the diversity in the human captions during training and results in models that tend to produce generic and commonly occurring captions from the training set, as we will show in Section 5.3.

We propose to address this issue by explicitly training the generator G to produce multiple diverse captions for an input image using the adversarial framework [13]. In adversarial framework, a generative model is trained by pairing it with adversarial discriminator which tries to distinguish samples produced by the generator from those of true data distribution. The generator is trained with the objective to fool the discriminator, which is optimal when G exactly matches the data distribution. This allows for training generative models without making any restrictive assumptions on the data distribution and using only samples from this distribution. The adversarial framework matches very well with our task. We want to generate sentence samples from the conditional distribution $P(s|x)$, but have only few samples from this distribution in form of human captions to learn from.

To enable adversarial training, we introduce a second network, $D(x, s)$, which takes as input an image x and a caption set $S_p = \{s_1, \ldots, s_p\}$ and classifies it as either real or fake. Providing a set of captions per image as input to the discriminator allows it to look at the diversity in the caption set while making the classification. The discriminator can penalize the generator for producing very similar or repeated captions for an image and thus encourage the generator to produce more diverse captions.

In our setting the discriminator is trained to classify the captions drawn from the reference captions set, $R(x) = \{r_0, \cdots, r_{k-1}\}$, as real while classifying the captions produced by the generator, $G(x)$, as fake. The generator G can now be trained using an adversarial objective, i.e. G is trained to fool the discriminator to classify $G(x)$ as real.

3.1. Caption generator

We use a near state-of-the art caption generator model based on [31]. It uses the standard encoder-decoder framework with two stages: the encoder model which extracts feature vectors from the input image and the decoder or the language model which translates these features into a word sequence.

Image features. Images are encoded using activations from a pre-trained convolutional neural network (CNN). It was found in [31] that it is beneficial to augment the CNN features with explicit object detection features for captioning. Following this, we also extract a small feature vector encoding the presence or absence of an object category and provide as input to the generator.

Language Model. Our decoder shown in Figure 3, is adopted from a Long-Short Term Memory (LSTM) based language model architecture presented in [31] for image captioning. It consists of a three-layered LSTM network with residual connections between the layers. The LSTM network takes two image features as input. The first image feature, referred to as x^i, is presented to the LSTM at only the time step 0 and shares the input matrix with the word vectors. The second image feature, x^p, is accessible by the LSTM at all time-steps and has its own input matrix.

The softmax layer at the output of the generator produces a probability distribution over the vocabulary at each step.

$$y_t = \text{LSTM}(w_{t-1}, x^p, y_{t-1}, c_{t-1}) \quad (1)$$

$$p(w_t|w_{t-1}, x) = \text{softmax}\left[\beta W_d * y_t\right], \quad (2)$$

where c_t is the LSTM cell state at time t and β is a scalar parameter which controls the peakyness of the distribution. We introduce this parameter β to control how large a hypothesis space the generator explores during adversarial training and found it crucial to have β in the range $(2.0, 3.0)$ to achieve stable GAN training. An additional uniform random noise vector z, is input to the LSTM in adversarial training to allow the generator to use the noise to produce diversity.

Discreteness Problem. To produce captions from the generator we could simply sample from this distribution $p(w_t|w_{t-1}, x)$, recursively feeding back the previously sampled word at each step, until we sample the END token. One can generate multiple sentences by sampling and pick the sentence with the highest probability as done in [11]. Alternatively we could also use greedy search approaches like beam-search. However, if we directly use these discrete samples as input to the discriminator, we cannot backpropagate through them as they are discontinuous.

Alternatives to overcome this are the reinforce rule/trick [37], using the softmax distribution, or using the Gumbel-Softmax approximation [17, 24].

Using policy gradient algorithms with the reinforce rule/trick [37] allows estimation of gradients through discrete samples [15, 2, 40, 21]. A drawback of this approach is the computational cost. The gradient estimation at intermediate stages of sentence generation needs an estimation of intermediate rewards. To get these intermediate rewards, we need to explore multiple possible completions using monte-carlo search techniques.

Another option is to input the softmax distribution to the discriminator instead of samples. We experimented with this, but we find that the discriminator finds it easy to distinguish between the softmax distribution produced by the generator and the sharp reference samples, and the GAN training fails.

The last option, which we rely on in this work, it to use a continuous relaxation of the samples encoded as one-hot vectors using the Gumbel-Softmax approximation proposed in [17] and [24]. It offers a technique to allow backpropagation through samples from a categorical distribution. The main benefit of this approach is that it plugs into the model as a differentiable node and does not need any additional steps to estimate the gradients.

The Gumbel-Softmax approximation consists of two steps. First Gumbel-Max trick is used to reparametrize sampling from a categorical distribution. So given a random variable r drawn from a categorical distribution parametrized by $\Theta = \theta_0, \cdots, \theta_{v-1}$, r can be expressed as:

$$r = \text{one_hot}\left[\arg\max_i \left(g_i + \log\theta_i\right)\right], \quad (3)$$

where g_i's are i.i.d. random variables from the standard Gumbel distribution. Next the argmax in Equation (3) is replaced with softmax to obtain a continuous relaxation of the discrete random variable r.

$$r' = \text{softmax}\left[\frac{g_i + \log\theta_i}{\tau}\right], \quad (4)$$

where τ is the temperature parameter which controls how close r' is to r, with $r' = r$ when $\tau = 0$.

We use straight-through variation of the gumbel-softmax approximation [17] at the output of our generator to sample words during the adversarial training. In the straight-through variation, sample r is used in the forward path and soft approximation r' is used in the backward path to allow backpropogation.

3.2. Discriminator model

The discriminator network, D takes as input a set of captions $S_p = \{s_1, \ldots, s_p\}$ and an image x and classifies S_p as either real or fake. Ideally, we would like D to do this classification based on two criteria: a) whether $s_i \in S_p$ describes the image correctly, b) whether the set S_p is diverse enough to match the diversity in human captions.

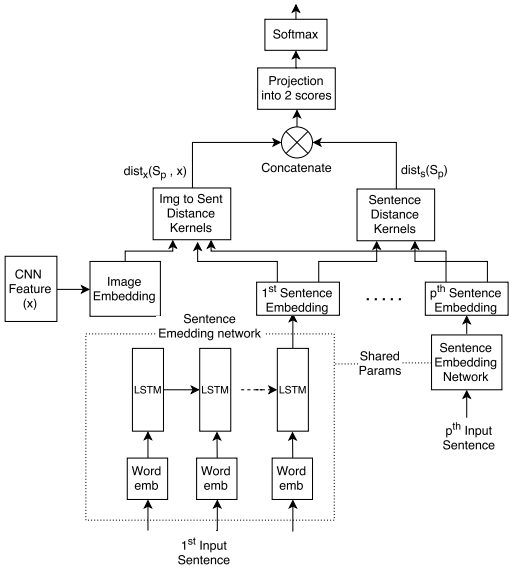

Figure 4: Block diagram of the Discriminator Network

To enable this, we use two separate distance measuring kernels in our discriminator architecture as shown in Figure 4. The first kernel computes the distances between the image x and each sentence in S_p. The second kernel computes the distances between the sentences in S_p. The architecture of these distance measuring kernels is based on the minibatch discriminator presented in [30]. However, in our case we only compute distances between multiple captions corresponding to the same image and not over the entire mini-batch.

First the input captions are encoded into a fixed size sentence embedding vector using an LSTM encoder to obtain vectors $f(s_i) \in \mathbb{R}^M$. The image feature vector is also embedded into an image embedding vector $f(x) \in \mathbb{R}^M$. The distances between $f(s_i), i \in \{1, \ldots, p\}$ are computed as

$$
\begin{aligned}
K_i &= T_s \cdot f(s_i) \\
c_l(s_i, s_j) &= exp\left(-\|K_{i,l} - K_{j,l}\|_{L_1}\right) \\
d_l(s_i) &= \sum_{j=1}^{p} c_l(s_i, s_j) \\
dist_s(S_p) &= [d_1(s_1), \cdots, d_O(s_1), \cdots d_O(s_p)] \in \mathbb{R}^{p \times O}
\end{aligned}
$$

where T_s is a $M \times N \times O$ dimensional tensor.

Distances between $f(s_i), i \in 1, \ldots, p$ and $f(x)$ are obtained with similar procedure as above, but using a different tensor T_x of dimensions $M \times N \times O$ to yield $dist_x(S_p, x) \in \mathbb{R}^{p \times O}$. These two distance vectors capture the two aspects we want our discriminator to focus on. $dist_x(S_p, x)$ captures how well S_p matches the image and $dist_x(S_p, x)$ captures the diversity in S_p. The two distance vectors are concatenated and multiplied with a output matrix followed by softmax to yield the discriminator output probability, $D(S_p, x)$, for S_p to be drawn from reference captions.

3.3. Adversarial Training

In adversarial training both the generator and the discriminator are trained alternatively for n_g and n_d steps respectively. The discriminator tries to classify $S_p^r \in R(x)$ as real and $S_p^g \in G(x)$ as fake. In addition to this, we found it important to also train the discriminator to classify few reference captions drawn from a random image as fake, i.e. $S_p^f \in R(y), y \neq x$. This forces the discriminator to learn to match images and captions, and not just rely on diversity statistics of the caption set. The complete loss function of the discriminator is defined by

$$
\begin{aligned}
L(D) = &-\log\left(D(S_p^r, x)\right) - \log\left(1 - D(S_p^g, x)\right) \\
&- \log\left(1 - D(S_p^f, x)\right) \quad (5)
\end{aligned}
$$

The training objective of the generator is to fool the discriminator into classifying $S_p^g \in G(x)$ as real. We found helpful to additionally use the feature matching loss [30]. This loss trains the generator to match activations induced by the generated and true data at some intermediate layer of the discriminator. In our case we use an l_2 loss to match the distances $dist_s(S_p)$ and $dist_x(S_p, x)$. The generator loss function is given by

$$
\begin{aligned}
L(G) = &-\log\left(D(S_p^g, x)\right) \\
&+\|\mathbb{E}\left[dist_s(S_p^g)\right] - \mathbb{E}\left[dist_s(S_p^r)\right]\|_2 \quad (6) \\
&+\|\mathbb{E}\left[dist_x(S_p^g, x)\right] - \mathbb{E}\left[dist_x(S_p^r, x)\right]\|_2,
\end{aligned}
$$

where the expectation is over a training mini-batch.

4. Experimental Setup

We conduct all our experiments on the MS-COCO dataset [5]. The training set consists of 83k images with five human captions each. All the results reported here, except in Section 5.4, use the standard publicly available test split of 5000 images [18]. Section 5.4 uses a validation split of 5000 images.

For image feature extraction, we use activations from *res5c* layer of the 152-layered *ResNet* [14] convolutional neural network (CNN) pre-trained on ImageNet. The input images are scaled to 448×448 dimensions for *ResNet* feature extraction. Additionally we use features from the VGG network [32] in our ablation study in Section 5.4. Following [31], we additionally extract 80-dimensional object detection features using a Faster Region-Based Convolutional Neural Network (RCNN) [28] trained on the 80 object categories in the COCO dataset. The CNN features are input to both the generator (at x_p) and the evaluator. Object detection features are input only to the generator at the x_i input and is used in all the generator models reported here.

4.1. Insights in Training the GAN

It is well known that GAN training is sensitive to hyperparameters [3]. This was observed in our experiments as well. Here we discuss some settings which helped stabilize the training of our models.

We found it necessary to pre-train the generator using standard maximum likelihood training. Without pre-training, the generator gets stuck producing incoherent sentences made of random word sequences. We also found pre-training the discriminator on classifying correct image-caption pairs against random image-caption pairs helpful to achieve stable GAN training.

We train the discriminator for 5 iterations for every update to the generator. Additionally we periodically monitor the classification accuracy of the discriminator and train it further if the accuracy drops below 75%. This helps prevent the generator from updating using a bad discriminator.

Without using the feature matching term in the generator loss, the GAN training was found to be unstable and needed a maximum likelihood update to the generator with every adversarial update. This was also reported in [21]. However using the feature matching loss removed the necessity for the ML update and we do not use ML update in any of the models reported here.

A good range of values for the Gumbel temperature was found to be $(0.1, 0.8)$. Beyond this range we found training to be unstable, but within this range the results were not sensitive to it. We use a fixed temperature setting of 0.5 in the experiments reported here. The softamax scaling factor, β in (2), is set to value 3.0 for training of all the adversarial models reported here. The sampling results are also with $\beta = 3.0$.

5. Results

We conduct experiments to evaluate our adversarial caption generator w.r.t. two aspects: how human-like the generated captions are and how accurately they describe the contents of the image. Using diversity statistics and word usage statistics as a proxy for measuring how closely the generated captions mirror the distribution of the human reference captions, we show that the adversarial model is more human-like than the baseline. Using human evaluation and automatic metrics we also show that the captions generated by the adversarial model performs similar to the baseline model in terms of correctness of the caption.

We refer to the baseline models as *Base* and the adversarial models as *Adv* in all tables. The extension *-bs* indicates captions obtained from beamsearch and *-samp* refer to captions obtained by sampling.

Method	Meteor	Spice
ATT-FCN [39]	0.243	–
MSM [38]	0.251	–
KWL [22]	0.266	**0.194**
Ours Base-bs	**0.272**	0.187
Ours Base-samp	0.265	0.186
Ours Adv-bs	0.239	0.167
Ours Adv-samp	0.236	0.166

Table 1: Meteor and Spice metrics comparing performance of baseline and adversarial models.

5.1. Measuring if captions are human-like

Diversity. We analyze n-gram usage statistics, compare vocabulary sizes and other diversity metrics presented below to understand and measure the gaps between human written captions and the automatic methods and show that the adversarial training helps bridge some of these gaps.

To measure diversity of the captions generated for the entire test set we use:

- *Vocabulary Size* - number of unique words used in all generated captions
- *% Novel Sentences* - percentage of generated captions not seen in the training set.

To measure diversity in a set of captions, S_p, corresponding to a single image we use:

- *div_1* - ratio of number of unique unigrams in S_p to number of words in S_p. Higher is more diverse.
- *div_2* - ratio of number of unique bigrams in S_p to number of words in S_p. Higher is more diverse.
- *mBleu* - Bleu score is computed between each caption in S_p against the rest. Mean of these p Bleu scores is the mBleu score. Lower values indicate more diversity.

Correctness. Just generating diverse captions is not useful if the captions do not correctly describe the content of an image. To measure the correctness of the generated captions we use two automatic evaluation metrics Meteor [7] and SPICE [1]. However since it is known that the automatic metrics do not correlate very well with human judgments of the correctness of a caption, we also report results from human evaluations comparing the baseline model to our adversarial model.

5.2. Comparing caption accuracy

Table 1 presents the comparison of our adversarial model to the baseline model. Both the baseline and the adversarial models use *ResNet* features. The beamsearch results are with beam size 5 and sampling results are with taking the best of 5 samples (in model's probability).

Table 1 also shows the metrics from some recent methods from the image captioning literature. The purpose of

Method	Spice					
	Color	Attribute	Object	Relation	Count	Size
Base-bs	**0.101**	**0.085**	0.345	0.049	0.025	0.034
Base-samp	0.059	0.069	**0.352**	**0.052**	0.032	0.033
Adv-bs	0.079	0.082	0.318	0.034	**0.080**	0.052
Adv-samp	0.078	0.082	0.316	0.033	0.076	**0.053**

Table 2: Comparing baseline and adversarial models in different categories of Spice metric.

Comparison	Adversarial - Better	Adversarial - Worse
Beamsearch	36.9	34.8
Sampling	35.7	33.2

Table 3: Human evaluation comparing adversarial model vs the baseline model on 482 random samples. Correctness of captions. With agreement of at least 3 out of 5 judges in %.

this comparison is to illustrate that we use a strong baseline and that our baseline model is competitive to recent published work, as seen from the Meteor and Spice metrics.

Comparing baseline and adversarial models in Table 1 the adversarial model does worse in-terms of meteor scores and overall spice metrics. When we look at Spice scores on individual categories shown in Table 2 we see that adversarial models excel at counting relative to the baseline and describing the size of an object right.

However, it is well known that automatic metrics do not always correlate with human judgments on correctness of a caption. A primary reason the adversarial models do poorly on automatic metrics is that they produce more unique sentences using much larger vocabulary and rarer n-grams, as will be shown in the following section. Thus, they are less likely to do well on metrics relying on n-gram matches.

To verify this claim, we conduct human evaluations comparing captions from the baseline and the adversarial model. Human evaluators are shown an image and a caption each from the two models and are asked "Judge which of the two sentences is a better description of the image (w.r.t. correctness and relevance)!". The choices were either of the two sentences or to report that they are the same. Results from this evaluation are presented in Table 3. We can see that both adversarial and baseline models perform similarly, with adversarial models doing slightly better. This shows that despite the poor performance in automatic evaluation metrics, the adversarial models produce captions that are similar, or perhaps even slightly better, in accuracy to the baseline model.

5.3. Comparing vocabulary statistics

To characterize how well the captions produced by the automatic methods match the statistics of the human written captions, we look at n-gram usage statistics in the generated captions. Specifically, we compute the ratio of the actual count of an n-gram in the caption set produced by a model to the expected n-gram count based on the training data.

Given that an n-gram occurred m times in the training set we can expect that it occurs $m * |\text{test-set}|/|\text{train-set}|$ times in the test set. However actual counts may vary depending on how different the test set is from the training set. We compute these ratios for reference captions in the test set to get an estimate of the expected variance of the count ratios.

The left side of Figure 5 shows the mean count ratios for uni-, bi- and tri-grams in the captions generated on test-set plotted against occurrence counts in the training set. Histogram of these ratios are shown on the right side.

Count ratios for the reference captions from the test-set are shown in green. We see that the n-gram counts match well with between the training and test set reference captions and the count ratios are distributed around 1.0 with a small variance.

The baseline model shows a clear bias towards more frequently occurring n-grams. It consistently overuses more frequent n-grams (ratio$>$1.0) from the training set and under-uses less frequent ones (ratio$<$1.0). This trend is seen in all the three plots, with more frequent tri-grams particularly prone to overuse. It can also be observed in the histogram plots of the count ratios, that the baseline model does a poor job of matching the statistics of the test set.

Our Adversarial model does a much better job in matching these statistics. The histogram of the uni-gram count ratios are clearly closer to that of test reference captions. It does not seem to be significantly overusing the popular words, but there is still a trend of under utilizing some of the rarer words. It is however clearly better than the baseline model in this aspect. The improvement is less pronounced with the bi- and tri-grams, but still present.

Another clear benefit from using the adversarial training is observed in terms of diversity in the captions produced by the model. The diversity in terms of both global statistics and per image diversity statistics is much higher in captions produced by the adversarial models compared to the baseline models. This result is presented in Table 4. We can see that the vocabulary size approximately doubles from 1085 in the baseline model to 2176 in the adversarial model using beamsearch. A similar trend is also seen comparing the sampling variants. As expected more diversity is achieved when sampling from the adversarial model instead of using beamsearch with vocabulary size increasing to 2671 in *Adv-samp*. The effect of this increased diversity can be in the qualitative examples shown in Figure 7. More qualitative samples are included in the supplementary material.

We can also see that the adversarial model learns to construct significantly more novel sentences compared to the baseline model with *Adv-bs* producing novel captions

Method	n	Div-1	Div-2	mBleu-4	Vocab-ulary	% Novel Sentences
Base-bs	1 of 5	–	–	–	756	34.18
	5 of 5	0.28	0.38	0.78	1085	44.27
Base-samp	1 of 5	–	–	–	839	52.04
	5 of 5	0.31	0.44	0.68	1460	55.24
Adv-bs	1 of 5	–	–	–	1508	68.62
	5 of 5	0.34	0.44	0.70	2176	72.53
Adv-samp	1 of 5	–	–	–	1616	73.92
	5 of 5	**0.41**	**0.55**	**0.51**	**2671**	**79.84**
Human captions	1 of 5	–	–	–	3347	92.80
	5 of 5	0.53	0.74	0.20	7253	95.05

Table 4: Diversity Statistics described in Section 5.1. Higher values correspond to more diversity in all except mBleu-4, where lower is better.

72.53% of the time compared to just 44.27% by the *beam-bs*. All three per-image diversity statistics also improve in the adversarial models indicating that they can produce a more diverse set of captions for any input image.

Table 4 also shows the diversity statistics on the reference captions on the test set. This illustrates that although adversarial models show considerable improvement over the baseline, there is still a gap in diversity statistics when compared to the human written captions, especially in vocabulary size.

Finally, Figure 6 plots the vocabulary size when considering words occurring at least k times, with varying k. This figure illustrates that the increase in vocabulary size in adversarial models is not obtained by just using some new words once but that it does indeed better mirror the vocabulary of the human captions in the test set.

5.4. Ablation Study

We conducted experiments to understand the importance of different components of our architecture on the validation set. The results are presented in Table 5. The baseline model for this experiment uses VGG [32] features as x_p input and is trained using maximum likelihood training and is shown in the first row of Table 5. The other four models use adversarial training.

We see that adversarial training with a discriminator evaluating a single caption does badly. Both the diversity and meteor score drop compared to the baseline. In this setting the generator can get away with producing one good caption for an input image as the discriminator is unable to penalize the lack of diversity in the generator.

However, as expected adversarial training using a discriminator evaluating 5 captions at a time does substantially better in terms of diversity.

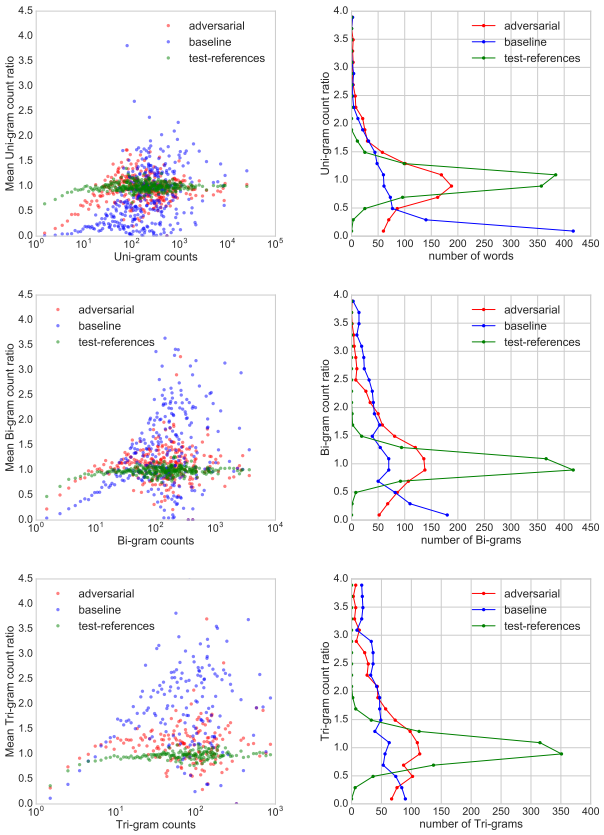

Figure 5: Comparison of n-gram count ratio statistics in captions generated on the test set by different models. Left side shows the mean n-gram count-ratios as a function of counts on training set. Right side shows the histogram of the count-ratios.

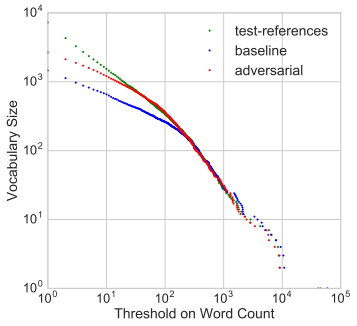

Figure 6: Vocabulary size as a function of word counts

Additionally using feature matching loss further improves diversity metrics and also slightly improves accuracy in terms of meteor score. Finally upgrading to the *ResNet* [14] gives a huge boost in terms of Meteor score and also slight increase in vocabulary size. With *ResNet* features, the

Adv-bs a group of friends enjoying a dinner at the restauarant | several cows in their pen at the farm | A dog is trying to get something out of the snow

Base-bs a group of people sitting around a wooden table | a herd of cattle standing next to each other | a couple of dogs that are in the snow

Figure 7: Some qualitative examples comparing comparing captions generated by the our model to the baseline model

Image Feature	Evalset size (p)	Feature Matching	Meteor	Div_2	Vocab. Size
VGG baseline			0.247	0.44	1367
VGG	1	No	0.179	0.40	812
VGG	5	No	0.197	0.52	1810
VGG	5	yes	0.207	**0.59**	2547
ResNet	5	yes	**0.236**	0.55	**2671**

Table 5: Performance comparison of various configurations of the adversarial caption generator on the validation set.

generator has access to richer visual information and can utilize this to produce diverse but still correct captions.

We also notice that the generator learns to ignore the noise input term. This is because there is sufficient stochasticity in the generation process due to sequential sampling of words and thus the generator doesn't need the additional noise input. Similar observation was reported in other conditional GAN architectures [16, 26]

6. Conclusions

In this work, we have presented an adversarial caption generator model which is explicitly trained to generate diverse captions for images. We achieve this by utilizing a discriminator network designed to promote diversity and use the adversarial learning framework to train our generator. Results show that our adversarial model produces captions which are diverse and match the statistics of human generated captions significantly better than the baseline model. The adversarial model also uses larger vocabulary and is able to produce significantly more novel captions. The increased diversity is achieved while preserving accuracy of the generated captions, as shown through a human evaluation.

Acknowledgements

This research was supported by the German Research Foundation (DFG CRC 1223) and by the Berkeley Artificial Intelligence Research (BAIR) Lab.

References

[1] P. Anderson, B. Fernando, M. Johnson, and S. Gould. Spice: Semantic propositional image caption evaluation. In *European Conference on Cmoputer Vision (ECCV)*, 2016. 6

[2] J. Andreas and D. Klein. Reasoning about pragmatics with neural listeners and speakers. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2016. 4

[3] M. Arjovsky and L. Bottou. Towards principled methods for training generative adversarial networks. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2017. 6

[4] S. Bengio, O. Vinyals, N. Jaitly, and N. Shazeer. Scheduled sampling for sequence prediction with recurrent neural networks. In *Advances in Neural Information Processing Systems (NIPS)*, 2015. 2

[5] X. Chen, T.-Y. L. Hao Fang, R. Vedantam, S. Gupta, P. Dollr, and C. L. Zitnick. Microsoft COCO captions: Data collection and evaluation server. *arXiv preprint arxiv:1504.00325*, 2015. 5

[6] COCO. Microsoft COCO Image Captioning Challenge. https://competitions.codalab.org/competitions/3221#results, 2017. 1

[7] M. Denkowski and A. Lavie. Meteor universal: Language specific translation evaluation for any target language. *ACL 2014*, 2014. 1, 6

[8] E. L. Denton, S. Chintala, R. Fergus, et al. Deep generative image models using a laplacian pyramid of adversarial networks. In *Advances in Neural Information Processing Systems (NIPS)*, 2015. 2, 3

[9] J. Devlin, H. Cheng, H. Fang, S. Gupta, L. Deng, X. He, G. Zweig, and M. Mitchell. Language models for image captioning: The quirks and what works. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*, 2015. 2

[10] J. Donahue, L. A. Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, K. Saenko, and T. Darrell. Long-term recurrent convolutional networks for visual recognition and description. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 2, 3

[11] J. Donahue, L. A. Hendricks, M. Rohrbach, S. Venugopalan, S. Guadarrama, K. Saenko, and T. Darrell. Long-term recurrent convolutional networks for visual recognition and description. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2016. 4

[12] A. Farhadi, M. Hejrati, M. Sadeghi, P. Young, C. Rashtchian, J. Hockenmaier, and D. Forsyth. Every picture tells a story: Generating sentences from images. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2010. 2

[13] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems (NIPS)*, 2014. 2, 3

[14] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 5, 8

[15] L. A. Hendricks, Z. Akata, M. Rohrbach, J. Donahue, B. Schiele, and T. Darrell. Generating visual explanations. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2016. 4

[16] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros. Image-to-image translation with conditional adversarial networks. *arXiv preprint arXiv:1611.07004*, 2016. 9

[17] E. Jang, S. Gu, and B. Poole. Categorical reparameterization with gumbel-softmax. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2016. 2, 3, 4

[18] A. Karpathy and L. Fei-Fei. Deep visual-semantic alignments for generating image descriptions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 2, 5

[19] G. Kulkarni, V. Premraj, V. Ordonez, S. Dhar, S. Li, Y. Choi, A. C. Berg, and T. L. Berg. Babytalk: Understanding and generating simple image descriptions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(12), 2013. 2

[20] J. Li, M. Galley, C. Brockett, J. Gao, and B. Dolan. A diversity-promoting objective function for neural conversation models. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, 2016. 2

[21] J. Li, W. Monroe, T. Shi, A. Ritter, and D. Jurafsky. Adversarial learning for neural dialogue generation. *arXiv preprint arXiv:1701.06547*, 2017. 2, 3, 4, 6

[22] J. Lu, C. Xiong, D. Parikh, and R. Socher. Knowing when to look: Adaptive attention via a visual sentinel for image captioning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 6

[23] P. Luc, C. Couprie, S. Chintala, and J. Verbeek. Semantic segmentation using adversarial networks. *arXiv preprint arXiv:1611.08408*, 2016. 2, 3

[24] C. J. Maddison, A. Mnih, and Y. W. Teh. The concrete distribution: A continuous relaxation of discrete random variables. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2016. 2, 4

[25] J. Mao, W. Xu, Y. Yang, J. Wang, Z. Huang, and A. Yuille. Deep captioning with multimodal recurrent neural networks (m-rnn). In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2015. 2

[26] M. Mathieu, C. Couprie, and Y. LeCun. Deep multi-scale video prediction beyond mean square error. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2016. 9

[27] A. Radford, L. Metz, and S. Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2016. 2, 3

[28] S. Ren, K. He, R. Girshick, and J. Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *Advances in Neural Information Processing Systems (NIPS)*, 2015. 5

[29] M. Rohrbach, W. Qiu, I. Titov, S. Thater, M. Pinkal, and B. Schiele. Translating video content to natural language descriptions. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2013. 2

[30] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen. Improved techniques for training gans. In *Advances in Neural Information Processing Systems (NIPS)*, 2016. 5

[31] R. Shetty, H. R-Tavakoli, and J. Laaksonen. Exploiting scene context for image captioning. In *ACMMM Vision and Language Integration Meets Multimedia Fusion Workshop*, 2016. 4, 5

[32] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2015. 5, 8

[33] R. Vedantam, C. Lawrence Zitnick, and D. Parikh. CIDEr: Consensus-based image description evaluation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 1

[34] A. K. Vijayakumar, M. Cogswell, R. R. Selvaraju, Q. Sun, S. Lee, D. Crandall, and D. Batra. Diverse beam search: Decoding diverse solutions from neural sequence models. *arXiv preprint arXiv:1610.02424*, 2016. 2

[35] O. Vinyals, A. Toshev, S. Bengio, and D. Erhan. Show and tell: A neural image caption generator. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 2, 3

[36] Z. Wang, F. Wu, W. Lu, J. Xiao, X. Li, Z. Zhang, and Y. Zhuang. Diverse image captioning via grouptalk. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, 2016. 2

[37] R. J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4), 1992. 3, 4

[38] T. Yao, Y. Pan, Y. Li, Z. Qiu, and T. Mei. Boosting image captioning with attributes. *arXiv preprint arXiv:1611.01646*, 2016. 6

[39] Q. You, H. Jin, Z. Wang, C. Fang, and J. Luo. Image captioning with semantic attention. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 6

[40] L. Yu, W. Zhang, J. Wang, and Y. Yu. SeqGAN: sequence generative adversarial nets with policy gradient. *Proceedings of the Conference on Artificial Intelligence (AAAI)*, 2016. 4

7. Supplementary Material

We present several qualitative examples to illustrate the strengths of our adversarially trained caption generator. All the examples are from the sampled versions of the adversarial (*adv-samp*) and the baseline (*base-samp*) models presented above. We show qualitative examples to highlight two main merits of the adversarial caption generator. First, we demonstrate diversity when sampling multiple captions for each image in Section 7.1. Next, we illustrate diversity across images in Section 7.2.

7.1. Illustrating diversity in captions for each image

To qualitatively demonstrate the diverse captions produced by the our adversarial model for each image, we visualize three captions produced by the adversarial and the baseline model for each input image. This is shown in figures 8 and 9. The captions are obtained by retaining the top three caption samples out of five (ranked by models' probability) from each model. Here bi-grams which are top-100 frequent bi-grams in the training set are highlighted in red (e.g., "a group" and "group of"). Additionally captions which are replicas from training set are marked with a '•' symbol. We can see that adversarial generator produces more diverse sets of captions for each image without over-using more frequent bi-grams and producing more novel sentences. For example, in the two figures, we see that the baseline model produces 22 captions (out of 45) which are copies from the training set, whereas the adversarial model does so only six times.

7.2. Illustrating diversity across images

Adversarial model produces diverse captions across different images containing similar content, whereas the baseline model tends to use common generic captions which is mostly correct but not descriptive. We quantify this by looking at the most frequently generated captions by the baseline model on the test set in Table 6. Note that here we consider only the most likely caption according to the model. Table 6 shows that the baseline model tends to repeatedly generate identical captions for different images. Compared to this, adversarial model is less prone to repeating generic captions, as seen in Table 7. This is visualized in Figures 10, 11 and 12. Here we show sets of five images for which the baseline model generates identical generic caption. The five images are picked from among the images corresponding to captions in Table 6, starting from the most frequent caption. Some entries are skipped, for example the caption in the last row, to avoid repeated concepts. While the baseline model produces a fixed generic caption for these images, we see that the adversarial model produces diverse captions and which is often more specific to the contents of image. For example, in the last row of Figure 10 we can see that the baseline model produces the generic caption "*a man riding skis down a snow covered slope*", whereas captions produced by the adversarial model include more image specific terms like "*jumping*", "*turn*", "*steep*" and "*corss country skiing*".

Sentence	**# baseline**	**# adversarial**
a man riding a wave on top of a surfboard	54	20
a bathroom with a toilet and a sink	44	1
a baseball player swinging a bat at a ball	37	2
a man riding skis down a snow covered slope	29	5
a man holding a tennis racquet on a tennis court	26	3
a bathroom with a sink and a mirror	25	3
a man riding a snowboard down a snow covered slope	24	4
a baseball player holding a bat on a field	22	1
a man riding a skateboard down a street	21	2
a bus is parked on the side of the road	20	0
...	...	...

Table 6: Most frequently repeated captions generated by the baseline model on the test set of 5000 images. The caption "a man riding a wave on top of a surfboard" is also the most frequently generated caption by the adversarial model, albeit less than half the times of the baseline model.

Sentence	**# adversarial**	**# baseline**
a man riding a wave on top of a surfboard	20	54
a skateboarder is attempting to do a trick	16	0
a female tennis player in action on the court	16	0
a living room filled with furniture and a flat screen tv	15	0
a bus that is sitting in the street	15	3
a long long train on a steel track	11	0
a close up of a sandwich on a plate	10	0
a baseball player swinging at a pitched ball	10	0
a bus that is driving down the street	9	1
a boat that is floating in the water	8	3
...	...	...

Table 7: Most frequently repeated captions generated by the adversarial model on the test set of 5000 images.

Adv-samp	• a red motorcycle parked on the side of the road • a motorcycle is parked on a city street a red motorcycle parked on a street in a city
Base-samp	• a man riding a motorcycle down a street • a person riding a motorcycle on a city street • a motorcycle is parked on the side of the road

Adv-samp	a motor cycle parked outside a building with people nearby • a motor bike parked in front of a building a row of bicycles parked outside of a building
Base-samp	• a group of people riding bikes down a street a group of people on a street with motorcycles a group of people on a street with motorcycles

Adv-samp	a motorcycle parked in front of a group of people a police officer on a motorcycle in front of a crowd of people a police officer on his motorcycle in front of a crowd
Base-samp	• a motorcycle parked on the side of a road a motorcycle is parked on the side of a road a motorcycle parked on the side of a road with a person walking by

Adv-samp	a skier is jumping over a snow covered hill a person on skis jumping over a hill a skier is in mid air after completing a jump
Base-samp	• a man riding skis down a snow covered slope • a man riding skis down a snow covered slope a person on a snowboard jumping over a ramp

Adv-samp	a group of people watching a skateboarder do stunts a group of skateboarders performing tricks at a skate park a group of skateboarders watch as others watch
Base-samp	• a man doing a trick on a skateboard at a skate park • a man riding a skateboard down a rail • a man doing a trick on a skateboard in a park

Adv-samp	• a stop sign with graffiti written on it a stop sign with a few stickers on it a stop sign has graffiti written all over it
Base-samp	a stop sign with a street sign on top of it a stop sign with a street sign on top of it • a stop sign with a street sign above it

Adv-samp	a man standing next to a mail truck a picture of people standing outside a business a couple of people standing by a bus
Base-samp	a group of people standing around a white truck a group of people standing around a white truck a group of people standing on a street next to a white truck

Adv-samp	a bouquet of flowers in a vase on a table a bouquet of flowers in a vase on a table a vase full of purple flowers sitting on a table
Base-samp	• a vase filled with flowers sitting on top of a table • a vase filled with flowers sitting on top of a table • a vase of flowers sitting on a table

Adv-samp	a plant in a vase on a wooden porch a small blue flower vase sitting on a wooden porch a large flower arrangement in a vase on a corner
Base-samp	• a vase with flowers in it sitting on a table • a vase of flowers sitting on a table • a vase with flowers in it on a table

Figure 8: Comparing 3 captions sampled from adversarial model to the baseline model. Bi-grams which are top-100 frequent bi-grams in the training set are highlighted in red. Captions which are replicas from training set are marked with a • .

Adv-samp	a long line of stairs leading to a church • a large cathedral filled with lots of pews a cathedral with stained glass windows and a few people	a large church with a very tall tower a large tall brick building with a clock on it a church steeple with a clock on its side	several stop signs in front of some buildings a stop sign in front of some graffiti writing several stop signs lined up in a row
Base-samp	a large building with a large window and a building a row of benches in front of a building a church with a large window and a large building	a large clock tower in a city a large clock tower in a city with a sky background • a large clock tower in the middle of a park	• a stop sign with graffiti on it a stop sign is shown with a lot of graffiti on it a stop sign and a stop sign in front of a building

Adv-samp	a family enjoying pizza at a restaurant party a group of friends enjoying pizza and drinking beer a group of kids enjoying pizza and drinking beer	a view of a city street at dusk a city street with many buildings and buildings a view of a city intersection in the evening	a white toilet sitting underneath a shower curtain • a white toilet sitting underneath a bathroom window a bathroom with a shower curtain open and a toilet in it
Base-samp	a group of people sitting at a table with pizza a group of people sitting at a table with pizza and drinks a group of people sitting at a table with pizza	a traffic light in a city with tall buildings a traffic light in a city with tall buildings • a traffic light and a street sign on a city street	• a bathroom with a toilet and a shower • a bathroom with a toilet and a shower a white toilet sitting next to a shower in a bathroom

Figure 9: Comparing 3 captions sampled from adversarial model to the baseline model.

Adv-samp: a surfer rides a large wave in the ocean | a surfer is falling off his board as he rides a wave | a person on a surfboard riding a wave | a man surfing on a surfboard in rough waters | a surfer rides a small wave in the ocean

Base-samp: . . . . . a man riding a wave on top of a surfboard . . . . .

Adv-samp: a bathroom with a walk in shower and a sink | a dirty bathroom with a broken toilet and sink | a view of a very nice looking rest room | a white toilet in a public restroom stall | a small bathroom has a broken toilet and a broken sink

Base-samp: . . . . . a bathroom with a toilet and a sink . . . . .

Adv-samp: a baseball player getting ready to swing at a pitch | a boy in a baseball uniform swinging a bat | a group of kids playing baseball on a field | a baseball game in progress with the batter upt to swing | a crowd watches a baseball game being played

Base-samp: . . . . . a baseball player swinging a bat at a ball . . . . .

Adv-samp: a person on skis jumping over a ramp | a skier is making a turn on a course | a cross country skier makes his way through the snow | a skier is headed down a steep slope | a person cross country skiing on a trail

Base-samp: . . . . . a man riding skis down a snow covered slope . . . . .

Figure 10: Illustrating diversity across images

Adv-samp: a tennis player gets ready to return a serve | two men dressed in costumes and holding tennis rackets | a tennis player hits the ball during a match | a male tennis player in action on the court | a man in white is about to serve a tennis ball

Base-samp: a man holding a tennis racquet on a tennis court

Adv-samp: a young boy riding a skateboard down a street | a skateboarder is attempting to do a trick | a boy wearing a helmet and knee pads riding a skateboard | a boy in white shirt doing a trick on a skateboard | a boy is skateboarding down a street in a neighborhood

Base-samp: a man riding a skateboard down a steet

Adv-samp: a dish with noodles and vegetables in it | a plate of food that has some fried eggs on it | a meal consisting of rice meat and vegetables | a close up of some meat and potatoes | a plate with some meat and vegetables on it

Base-samp: a plate of food with meat and vegetables

Adv-samp: a group of people standing around a shop | a group of young people standing around talking on cell phones | a group of soldiers stand in front of microphones | a couple of women standing next to a man in front of a store | a group of people posing for a photo in formal wear

Base-samp: a group of people standing around a table

Figure 11: Illustrating diversity across images

Adv-samp: a group of men sitting around a meeting room | a group of people sitting at a bar drinking wine | a group of friends enjoying lunch outdoors at a outdoor event | a group of people sitting at tables outside | a couple of men that are working on laptops

Base-samp: . . . . . a group of people sitting around a table . . . . . .

Adv-samp: a laptop and a desktop computer sit on a desk | a person is working on a computer screen | a cup of coffee sitting next to a laptop | a laptop computer sitting on top of a desk next to a desktop computer | a picture of a computer on a desk

Base-samp: . . . . . a laptop computer sitting on top of a desk . . . . . .

Figure 12: Illustrating diversity across images